REMARKS

Reconsideration of this application is respectfully requested in view of the foregoing amendment and the following remarks.

Claims 1-8, 12-19, 23-30, 33-44, and 47-50 were pending in this application. An Amendment After Final was filed on April 12, 2005, but was not entered. By the current Amendment, claims 1-8 and 12-19 have been cancelled without prejudice or disclaimer to the subject matter thereof, and claims 23, 33-37, and 47-50 have been amended. Accordingly, claims 23-30, 33-44, and 47-50 will be pending herein upon entry of this Amendment. For at least the reasons stated below, Applicants respectfully submit that all claims pending in this application are in condition for allowance.

In the Office Action mailed April 7, 2004, claims 1, 2, 5, 7, 8, 12, 13, 16, 18, 19, 23, 24, 27, 29, 30, 37, 38, 41, 43, and 44 were rejected under 35 USC § 102(e) as being anticipated by Munson (US Patent No. 6,681, 331) and claims 3, 4, 6, 14, 15, 17, 25, 26, 28, 39, 40, and 42 were rejected under 35 USC § 103(a) as being unpatentable over Munson, in view of Rowland (US Patent No. 6,405,318), which are addressed below. To the extent these rejections might still be applied to claims presently pending in this application, they are respectfully traversed.

Regarding the anticipation rejection of claims 23 and 37, as previously stated, Applicants note that Munson generally teaches a manner of anomaly detection by comparing events generated on a computer to events deemed to be normal and determining whether the generated event is normal or abnormal, but the specific manner that Munson teaches to determine whether an event is normal or abnormal is quite different from the present invention. The system

described in Munson utilizes multinomial distributions to determine abnormality. The present system employs neural networks that have been previously trained to identify normal behavior. The detected behavior is then analyzed, often in real-time, to determine if that particular data stream is normal.

These trained neural networks are trained during a training phase as described, for example, starting at the bottom of page 15. During the training phase, the neural networks learn the types of behaviors that are considered normal as well as the types that are considered intrusive. The training data is collected and labeled by a system administrator accordingly so that the neural network knows how to classify various behaviors. Once the training phase is completed, a new set of test data is used to test the trained neural networks to see how they perform.

During the testing phase, the neural network may encounter new behavior that was not specifically designated as normal or intrusive during the training phase. Because neural networks are employed, however, the intrusion detection system has the ability to classify new behavior without the need for a system administrator to specifically designate the behavior as such. Accordingly, the trained neural networks may be judged according to performance during the testing phase, to see which one performed best. The trained neural network that has performed the best given the specific type of application monitoring to be performed may then be selected and implemented.

Munson, on the other hand, can detect only based upon criteria specifically laid out by a system administrator. Munson even specifically recognizes this limitation in its functionality in

column 6, lines 46+ where it describes that only known intrusion events will raise a level 2 alarm. Anything classified as new will raise a level 1 alarm that has to be reviewed by a system administrator or an undescribed AI engine. Munson goes on to describe at line 66, that human pattern recognition surpasses any available software and therefore the use of a visualizer monitored by a human system administrator is the preferable manner of implementation. Contrary to Munson's teaching, the present invention eliminates the need for further monitoring by the system administrator and it is the trained neural network that is capable of determining whether new behavior is indeed normal or intrusive.

At least claims 23 and 37 (and by incorporation each of their respective dependent claims) include specific reference to the above-referenced "trained neural networks", which are not employed, or even discussed in Munson. Examiner's indication of this element being taught at col. 7, lines 17-20, is not equivalent (nor does it even mention) neural networks that are "trained to identify a pre-determined behavior pattern for a corresponding one of the plurality of applications." Nor is data collected via an application profile "sequentially input into a corresponding one of the plurality of trained neural networks" in order to obtain "a behavior indicator for each of the plurality of data strings in the application profile." In fact, the analysis performed by Munson relates to multinomial statistical analysis of detected events versus known normal events and does not involve the use of trained neural networks as described in claims 23 and 37. Although Applicants continue to disagree with the Examiner's characterization of Munson as employing neural networks at all, Applicants have amended the claims to clarify this training aspect. In order for a rejection to be maintained under 35 USC § 102, each and every

limitation of the claims must be present in the cited reference. Clearly at least the use of trained neural networks for intrusion detection as claimed is not taught by Munson.

Further, Applicants claimed "application profiles" are not equivalent to the operational profiles utilized by Munson, as suggested by the Examiner. Munson defines his operational profiles at col. 9, lines 20-25 as "[t]he set of unconditional probabilities of each of the operations in O being executed by the user." Munson then bases its intrusion detection upon these operational profiles by comparing what is occurring in the operational profiles to pre-existing intrusion profiles data (col. 4, lines 26-65). This type of comparison is not what is claimed in claims 23 or 37. Applicants' application profiles "comprise a plurality of application data for a corresponding one of the plurality of applications," and have nothing to do with unconditional probabilities. In fact, the present approach forgoes the computation of probabilities, which often provides for better results than explicitly computing probabilities as is performed in Munson. Next, "a behavior indicator for each of the plurality of data strings in the application profile" is output and "if the behavior indicator meets a pre-determined criteria, a counter is incremented." This is not the same analysis performed in Munson and the Examiner has pointed to no teaching in Munson that could be considered equivalent.

Regarding the rejection based upon the combination of Munson with Bergman et al., such a combination is essentially infeasible. The nodes referred to in Bergman are nodes of the communications network itself, not of a trained neural network (which as noted above is not even disclosed by Munson). The nodes recited in claims 33-35 and 47-49 reside in the plurality of trained neural networks, and are utilized to monitor whatever application is being run.

Further, backpropogation is a learning algorithm utilized in neural networks; there simply is no such thing as backpropogation of a communications network. Backpropogation occurs before the neural network is deployed. It is impossible for the backpropogation of a neural network *before* deployment to contain evidence of activity deemed intrusive *after* the network is deployed. Accordingly any combination of Munson with Bergman for the purpose described by the Examiner must fail.

As to Rowland, because this reference is applied only to dependent claims and fails to supply any of the above-cited deficiencies with respect to the Munson reference, Applicants assert that the combination of Rowland and Munson still does not teach each element of the independent claims at issue.

In summation, Applicants do not assert to have invented intrusion detection. Applicants do however, believe the present application represents a novel manner of intrusion detection involving training neural networks during a training phase. In certain embodiments, the trained neural networks are then selected for use in application monitoring based upon performance during a testing phase, and then are used for monitoring applications utilizing specifically selected trained neural networks. There simply is no discussion in the cited references of the type of training of neural networks for use in application monitoring as claimed in claims 23 or 37. While the cited references my involve a type of intrusion detection, it is not the type specifically claimed in either of independent claims 23 or 37 or their associated dependent claims.

Applicants have noted numerous deficiencies in the application of Munson, both alone and in combination with other references. Applicants maintain that the Office has failed in its burden to show that each and every recitation of the claims is present in the cited references. While Applicant still maintains that the previous office action suffered deficiencies in application of the cited art and that many claim recitations were not present in the cited references, the present amendments are believed to clarify these positions and further highlight the deficiencies of the cited references *vis-à-vis* the pending claims. Accordingly, because each and every element recited in each of the independent claims is not present in either the Munson, Bergman, Rowland references, or any combination thereof, Applicants assert that claims 23 and 37 along with their respective dependent claims are in condition for allowance. Should the Examiner have any questions or determine that any further action is desirable to place this application in even better condition for issue, the Examiner is encouraged to telephone applicants' undersigned representative at the number listed below.

PILLSBURY WINTHROP SHAW PITTMAN LLP
1650 Tysons Boulevard
McLean, VA 22102
Tel: 703-770-7900

Respectfully submitted

ANUP K. GHOSH, ET AL.

Date: June 13, 2005

By: 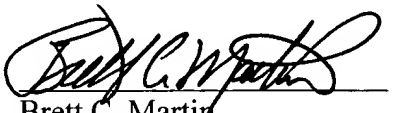
Brett C. Martin
Registration No. 52,000

BCM/dkp

Customer No. 28970